



cyber reboot[®]

Finding Aram Fingal:

Using Poseidon and Faucet to Identify and Monitor Insider Threats.



Who is Aram Fingal?



Hapless engineer lost in the network?

Poseidon moves us closer to “intrusion tolerance,” where we start with the assumption that a network has already been compromised. Can we use ML and SDN to identify adversaries before the objective is reached?

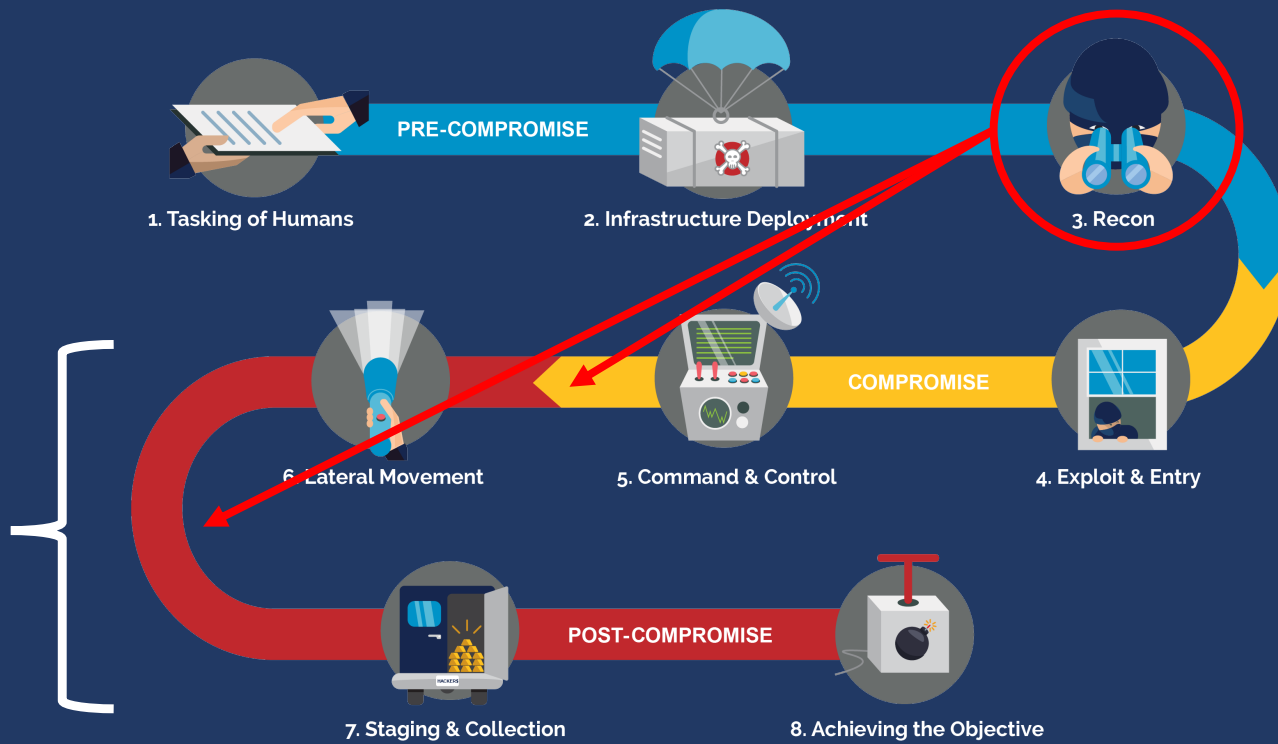
Or



Devious threat plotting nefarious action?

Deception techniques present the potential to increase complexity for adversaries, assist with detection, and automate some analysis. Can we achieve any of these goals without making more work for the defenders?

Attacker Workflow



Efforts in Post-Compromise
Poseidon
Deception
Telemetry Enrichment

How are we going to find him?

Moving toward better intrusion tolerance by hindering an attacker's ability to proceed through an attack chain



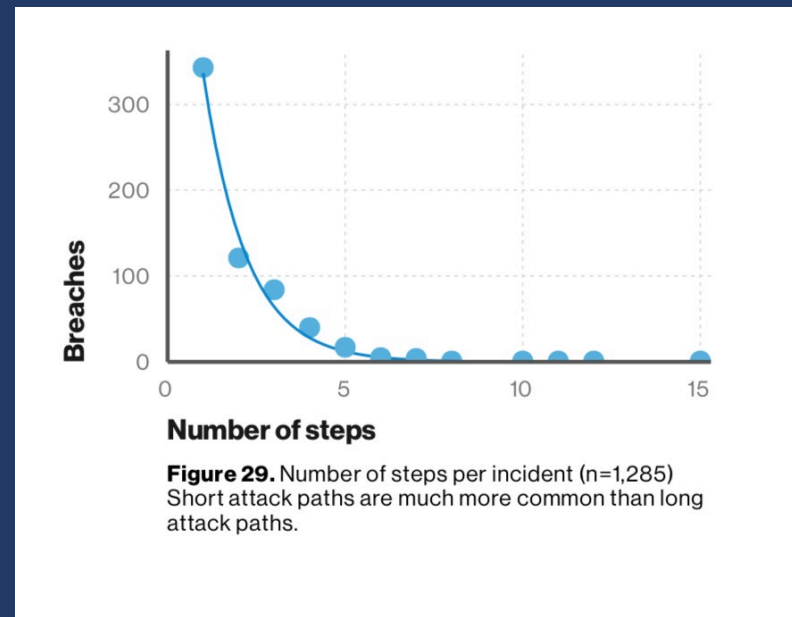
We are going to find him by keeping him in the “Recon” phase and forcing him to return to it more often. How are we going to do that?

Deception!

- Not Just Honeypots!
- Deception provides one class of sensors
- Federate data from multiple types of sensors to create stronger analytics
- Algorithmically determine the response based on data

Why are we doing this? And why now?

- Current techniques aren't working
 - Breach volume is up
 - Dwell times continue to rise
- Defense being a human only operation is unsustainable
 - Personnel shortage to begin with
 - High Burnout rate
 - Inherently reactive
- Need to move towards intrusion tolerant Networks
- SDN and Orchestration platforms provide a new opportunity



*Source: Verizon 2019 Data Breach Incident Report - <https://enterprise.verizon.com/resources/reports/dbir/>

Why Deception and SDN complement each other

TLDR: API's

Longer More Nuanced Answer:

Because an SDN can be aware of the traffic passing through it and can respond to events taking place on the network in near real time. So instead of a human operator needing to:

1. Be alerted to some kind of incident or anomaly
2. Determine the source
3. Verify What is going on
4. Respond appropriately

The human only needs to deal with 2 things:

1. Is this a false positive?
2. If not, was the automated action sufficient

This puts a machine doing what machines are good at (reacting to stimuli by rote) and puts the human in the position of doing what humans do well: creative thinking and analysis. More critically, this changes operator thinking patterns from an incident response mindset (i.e. where is the evidence) to a proactive, engagement planning mindset.

Rules.yaml and ACL setup

Everybody loves YAML!

```
1 include:
2   - acls.yaml
3
4 rules:
5   rule-name-1:
6     - rule:
7       device_key: os
8       value: Mac
9       acls: [office-vlan-protect, foo]
10
11     - rule:
12       device_key: role
13       value: Printer
14       min_confidence: 50
15       acls: [no-external]
16
17   rule-name-2:
18     - rule:
19       device_key: behavior
20       value: abnormal
21       acls: [no-internal]
22
23 no-internal:
24   - rule:
25     acls: [no-internal]
```

Can be applied based on a ruleset

Or specifically set

ACLs can be defined either in FAUCET's config or in included in the ruleset file. If the included ACL's do not exist in FAUCET's config, Poseidon will copy them over before applying them. ACL's can be applied automatically by ruleset, or set by sending a message on the message bus



The Honeyseidon script

Even more YAML!

```
whitelist: []
blacklist: []
logfile: '/var/log/honeyposeidon.log'
rabbit:
  FA_RABBIT_ENABLED: True
  FA_RABBIT_HOST: localhost
  FA_RABBIT_PORT: 5672
  FA_RABBIT_EXCHANGE: topic_recs
  FA_RABBIT_EXCHANGE_TYPE: topic
  FA_RABBIT_ROUTING_KEY: FAUCET.Event
ips:
  '127.0.0.1':
    ports:
      80:
        protocol: 'tcp'
        response_script: 'action.py'
      443:
        protocol: 'tcp'
        response_script: 'action.py'
      31337:
        protocol: 'tcp'
        #response_script: 'action.py'
```

There are 2 important parts to the script. The first part is a yamll file that defines configuration specifically:

- IPs
- Ports
- Protocols

To listen on and a script to execute upon receiving a connection on the specified IP/port/protocol combo. Note that each instance can have individual response scripts.

Secondly, the response scripts are simply a set of user defined python scripts that receive the source and destination IPs and ports as parameters.

```
def respond(src_ip, src_port, dest_ip, dest_port):
    base_uri = "http://localhost:5000/v1/"
    network_resource = "network"
    network_uri = urljoin(base_uri, network_resource)
    session = requests.Session()
    response = session.get(network_uri, verify=False)
```


Side Effect: Interfering with banner grabbing



In addition to a response script, response text can be defined as well. You could use this as a way to politely inform someone that they are unwelcome to connect to that port on that machine.

Or you could use it to reinforce your deception:

- Send an IIS banner on port 80 and an Apache banner on port 443
- Send an SSH banner on port 139 and 445
- Telnet banners on 15 different ports
- Send back non-sensical ascii art

The entire idea is to confuse your attackers and drive them nuts. Have fun with it. Let your imagination run wild.

Scenario 1: Abnormal Dev Workstation

A Dev Workstation deemed to be behaving abnormally is trying to connect to SSH on a webserver it should not be. We are going to respond by:

- Applying an ACL cutting off connection internally to disrupt pivoting and lateral movement activity
- Monitoring to determine if the activity returns to normal

Scenario 1: Abnormal Dev Workstation



Refining the Techniques

Overall Goal: Force attackers into treating the network as an “unreliable narrator”

New Faucet features we could Use:

- Coprocessing – can we shape/manipulate traffic in interesting or useful ways? Can we selectively break Idempotence for APIs, what are the results/consequences?
- Throttling– How will attackers react to variability of traffic speeds

Other things to add to the script:

- Dockerization – it would be cool to be able to just run this from a container I just have to figure out how to handle the port mappings. Docker-compose? Lazy-docker?
- Inconsistent Service behavior -

Key Questions:

- Effectiveness – run adversarial tests and see if it really is forcing attackers to take more steps.
- Operational Impacts – Does it hamper adversaries more than it does legitimate users.

Questions?



<https://github.com/cyberreboot/>